

(12) **United States Patent**
Clark et al.

(10) **Patent No.:** **US 9,298,511 B2**
(45) **Date of Patent:** ***Mar. 29, 2016**

(54) **RESOLVING DEPLOYMENT CONFLICTS IN HETEROGENEOUS ENVIRONMENTS**

(56) **References Cited**

U.S. PATENT DOCUMENTS

(71) Applicant: **International Business Machines Corporation**, Armonk, NY (US)

(72) Inventors: **Adam T. Clark**, Mantorville, MN (US);
Michael T. Kalmbach, Elgin, MN (US);
John E. Petri, St. Charles, MN (US);
Kevin Wendzel, Rochester, MN (US)

(73) Assignee: **International Business Machines Corporation**, Armonk, NY (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 298 days.

This patent is subject to a terminal disclaimer.

6,990,662 B2	1/2006	Messer et al.	
7,617,498 B1	11/2009	Lumsden	
8,090,837 B2	1/2012	Shin et al.	
8,584,128 B1	11/2013	Don et al.	
8,631,099 B2 *	1/2014	Morgan	709/220
8,972,956 B2	3/2015	Brownlow et al.	
2004/0158602 A1	8/2004	Broberg	
2005/0033767 A1	2/2005	Kamentz et al.	
2005/0251802 A1	11/2005	Bozek et al.	
2005/0262509 A1	11/2005	Kelly	
2006/0136908 A1 *	6/2006	Gebhart et al.	717/177
2009/0049443 A1	2/2009	Powers et al.	
2009/0288083 A1	11/2009	Ciano et al.	
2009/0293056 A1	11/2009	Ferris	
2010/0223378 A1	9/2010	Wei	

(Continued)

OTHER PUBLICATIONS

U.S. Application entitled Detecting Deployment Conflicts in Heterogeneous Environments, filed Mar. 15, 2013.

(Continued)

(21) Appl. No.: **13/840,703**

(22) Filed: **Mar. 15, 2013**

(65) **Prior Publication Data**

US 2014/0282573 A1 Sep. 18, 2014

(51) **Int. Cl.**

G06F 9/50 (2006.01)

G06F 9/45 (2006.01)

G06F 9/48 (2006.01)

(52) **U.S. Cl.**

CPC **G06F 9/5044** (2013.01); **G06F 8/443** (2013.01); **G06F 9/4806** (2013.01); **G06F 9/4881** (2013.01); **G06F 9/50** (2013.01); **G06F 9/5027** (2013.01)

(58) **Field of Classification Search**

CPC ... G06F 9/5044; G06F 9/4806; G06F 9/5027; G06F 9/50; G06F 9/4881; G06F 8/443

See application file for complete search history.

Primary Examiner — Charles E Anya

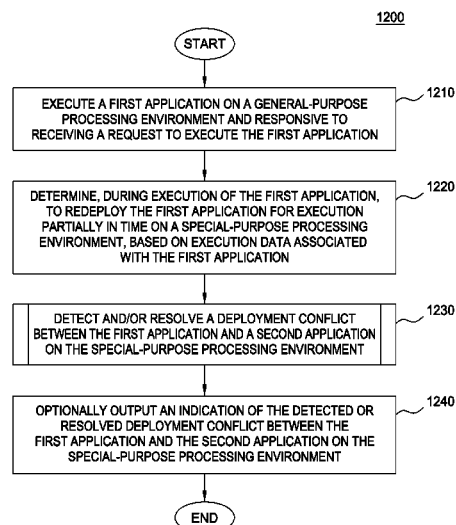
(74) *Attorney, Agent, or Firm* — Patterson & Sheridan, LLP

(57)

ABSTRACT

Techniques are disclosed for managing deployment conflicts between applications executing in one or more processing environments. A first application is executed in a first processing environment and responsive to a request to execute the first application. During execution of the first application, a determination is made to redeploy the first application for execution partially in time on a second processing environment providing a higher capability than the first processing environment in terms of at least a first resource type. A deployment conflict is resolved between the first application and at least a second application.

20 Claims, 14 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

2010/0287549 A1 11/2010 Neft
 2010/0318658 A1 12/2010 Zorn et al.
 2011/0131329 A1 6/2011 Kaplinger et al.
 2011/0131499 A1 6/2011 Ferris et al.
 2011/0213875 A1 9/2011 Ferris et al.
 2011/0296025 A1* 12/2011 Lieblich G06F 9/5027
 709/226
 2012/0023500 A1 1/2012 Karstens
 2012/0089726 A1* 4/2012 Doddavula 709/224
 2012/0102452 A1 4/2012 Guillou et al.
 2012/0102498 A1 4/2012 Subramanya et al.
 2012/0149464 A1 6/2012 Bone et al.
 2012/0192197 A1 7/2012 Doyle et al.
 2012/0297163 A1* 11/2012 Breternitz et al. 712/22
 2013/0031559 A1* 1/2013 Alicherry G06F 9/5077
 718/104
 2013/0055280 A1* 2/2013 Kruglick G06F 9/5027
 718/104
 2013/0097597 A1 4/2013 Gibson et al.
 2013/0232498 A1* 9/2013 Mangtani G06F 9/5072
 718/104

2013/0311628 A1* 11/2013 Kruglick 709/223
 2013/0326507 A1* 12/2013 McGrath et al. 718/1
 2013/0332588 A1* 12/2013 Maytal G06F 11/3428
 709/223

OTHER PUBLICATIONS

Sou Mya Mahadevan, "Performance Analysis of Offloading Application-Layer Tasks to Network Processors," Thesis, University of Massachusetts Amherst, Sep. 2007, Amherst, Massachusetts, USA, pp. 1-55.

Robert Filepp et al., "Image Selection as a Service for Cloud Computing Environments," 2010 IEEE International Conference on Service-Oriented Computing and Applications (SOCA), Dec. 2010, 8 pages.

U.S. Appl. No. 14/568,790, entitled "Resolving Deployment Conflicts in Heterogeneous Environments", filed Dec. 12, 2014.

U.S. Appl. No. 14/568,618, entitled "Detecting Deployment Conflicts in Heterogeneous Environments", filed Dec. 12, 2014.

* cited by examiner




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5

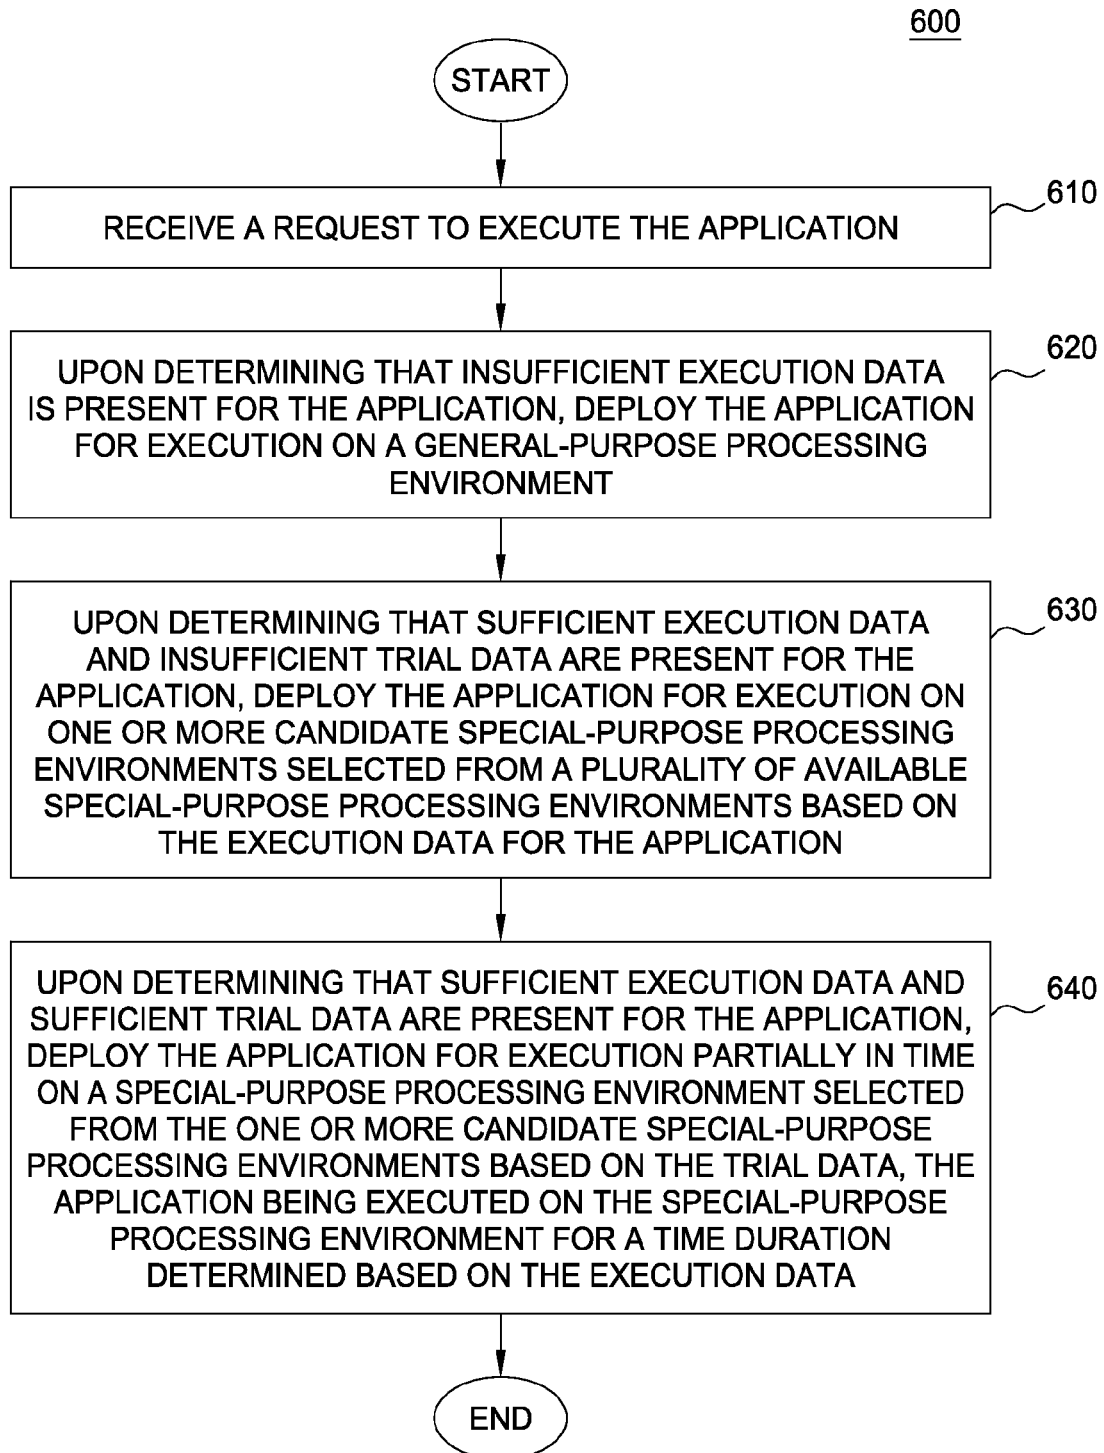


FIG. 6

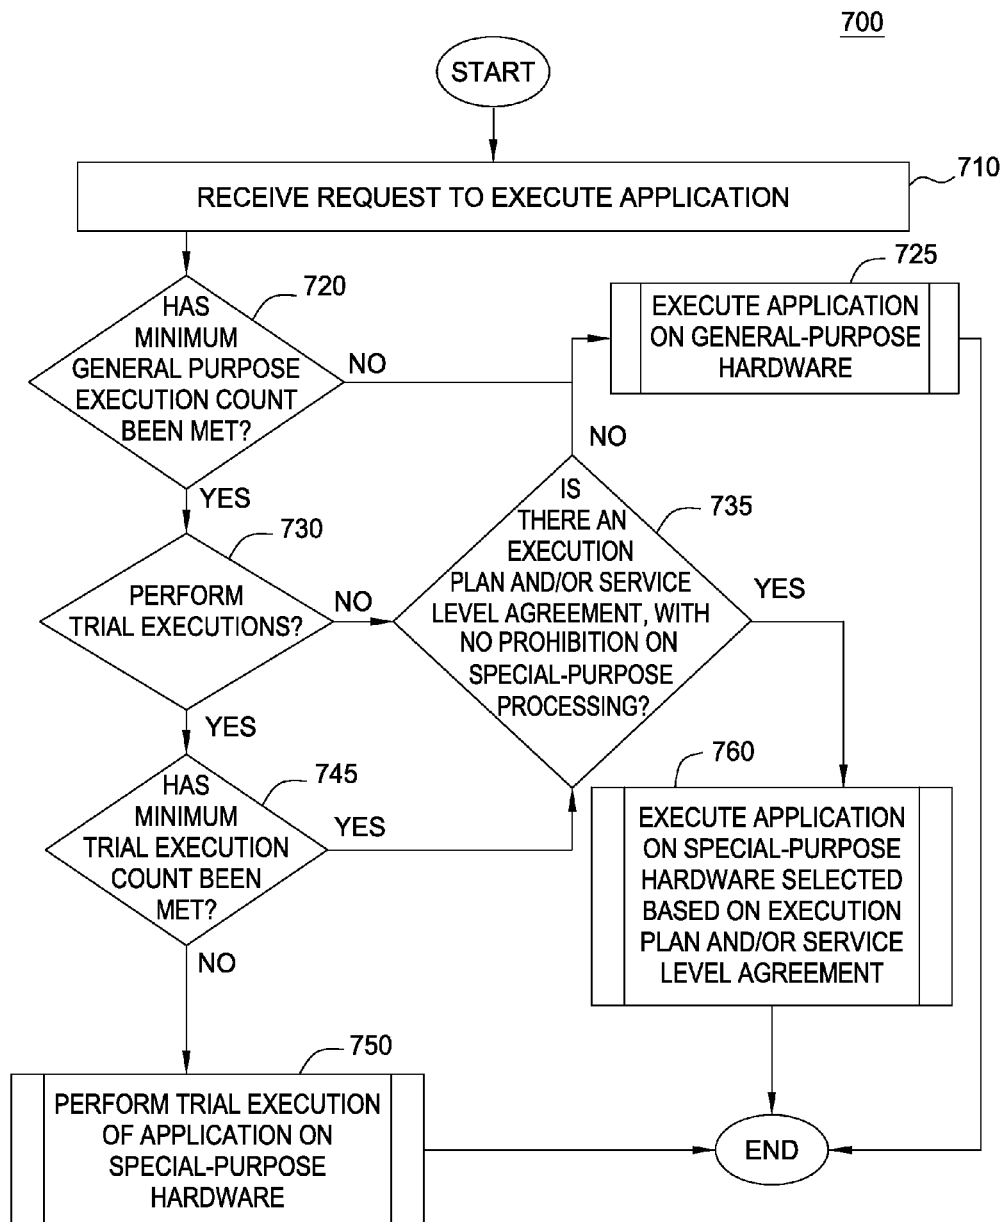


FIG. 7


800

FIG. 8

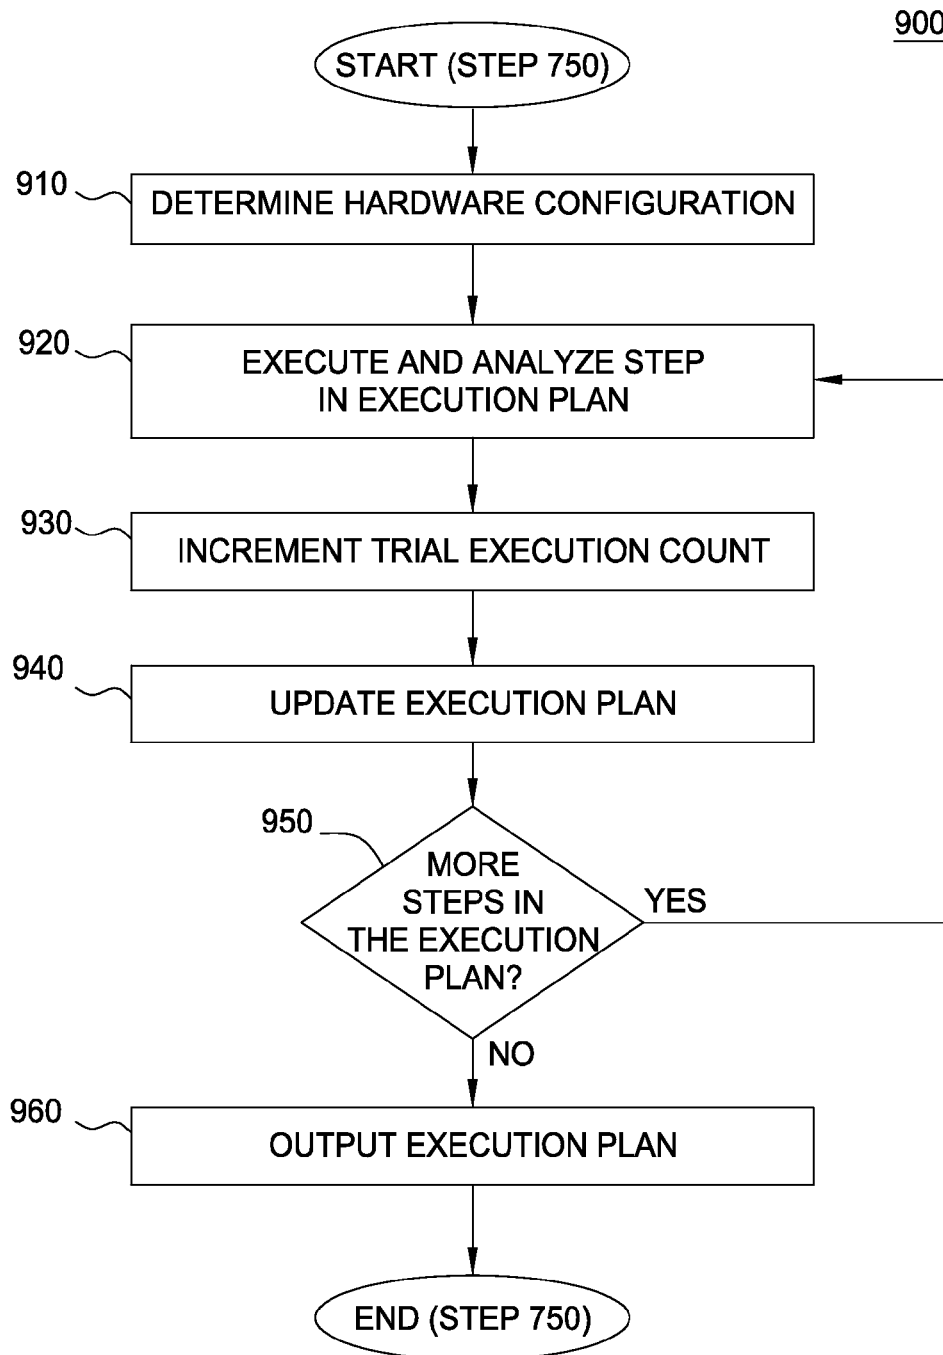


FIG. 9

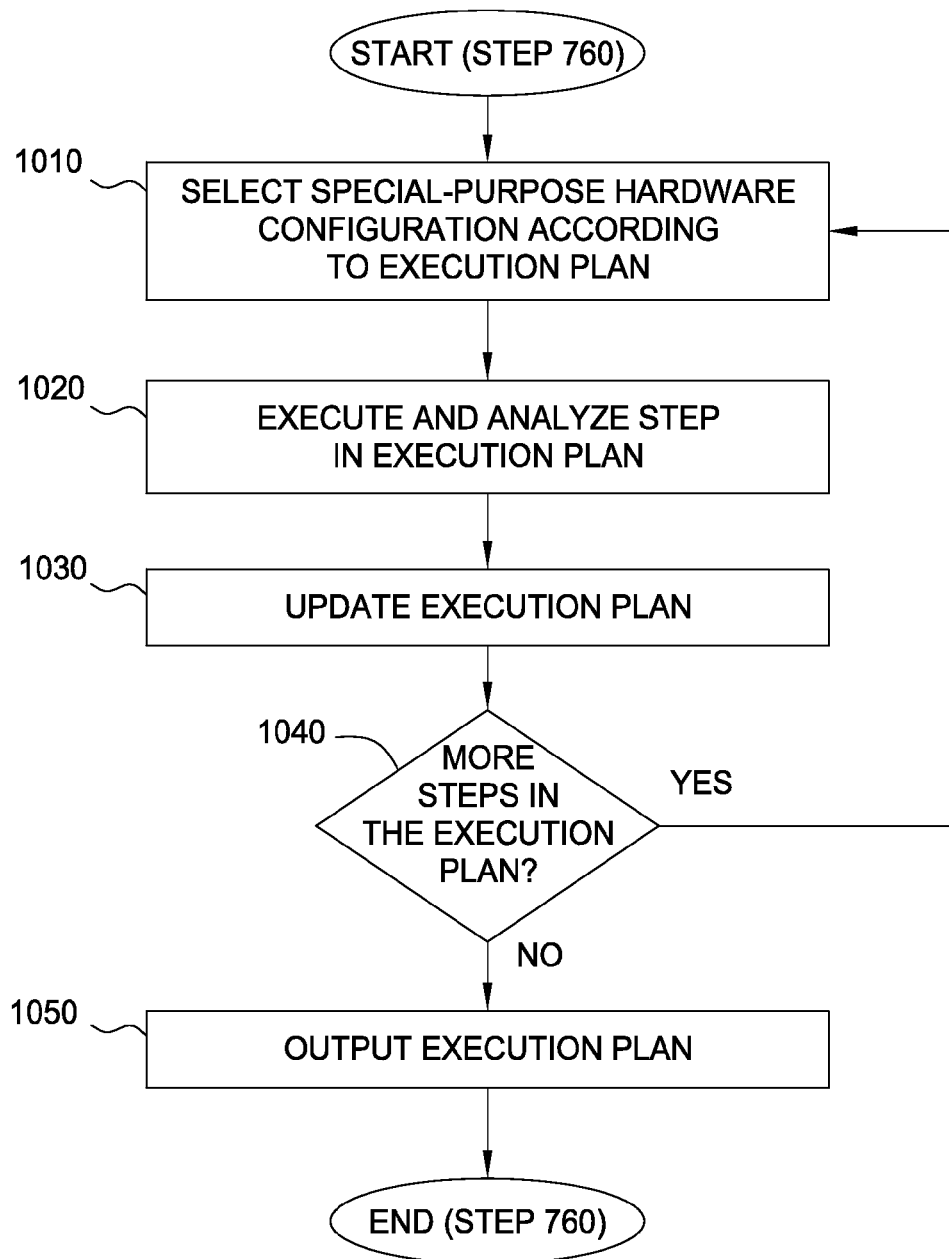
1000

FIG. 10

1100

1102			1104			1106		
EXECUTION PLAN A (VIDEO)			EXECUTION PLAN B (GAME)			EXECUTION PLAN C (SHOPPING APPLICATION)		
I/O BOUND (BUFFERING DATA, 10 SECONDS)			I/O BOUND (20 SECONDS)			I/O BOUND (1 SECOND)		
GRAPHICS BOUND (PLAYING MOVIE, 60 SECONDS)			USER INPUT (VARIES, TYPICALLY 10 TO 60 SECONDS)			USER INPUT (VARIES, TYPICALLY 30 TO 90 SECONDS)		
I/O BOUND (RELATED VIDEOS, 2 SECONDS)			GRAPHICS BOUND (LONG TIME)			GENERAL PURPOSE PROCESSING (10-20 SECONDS)		

FIG. 11

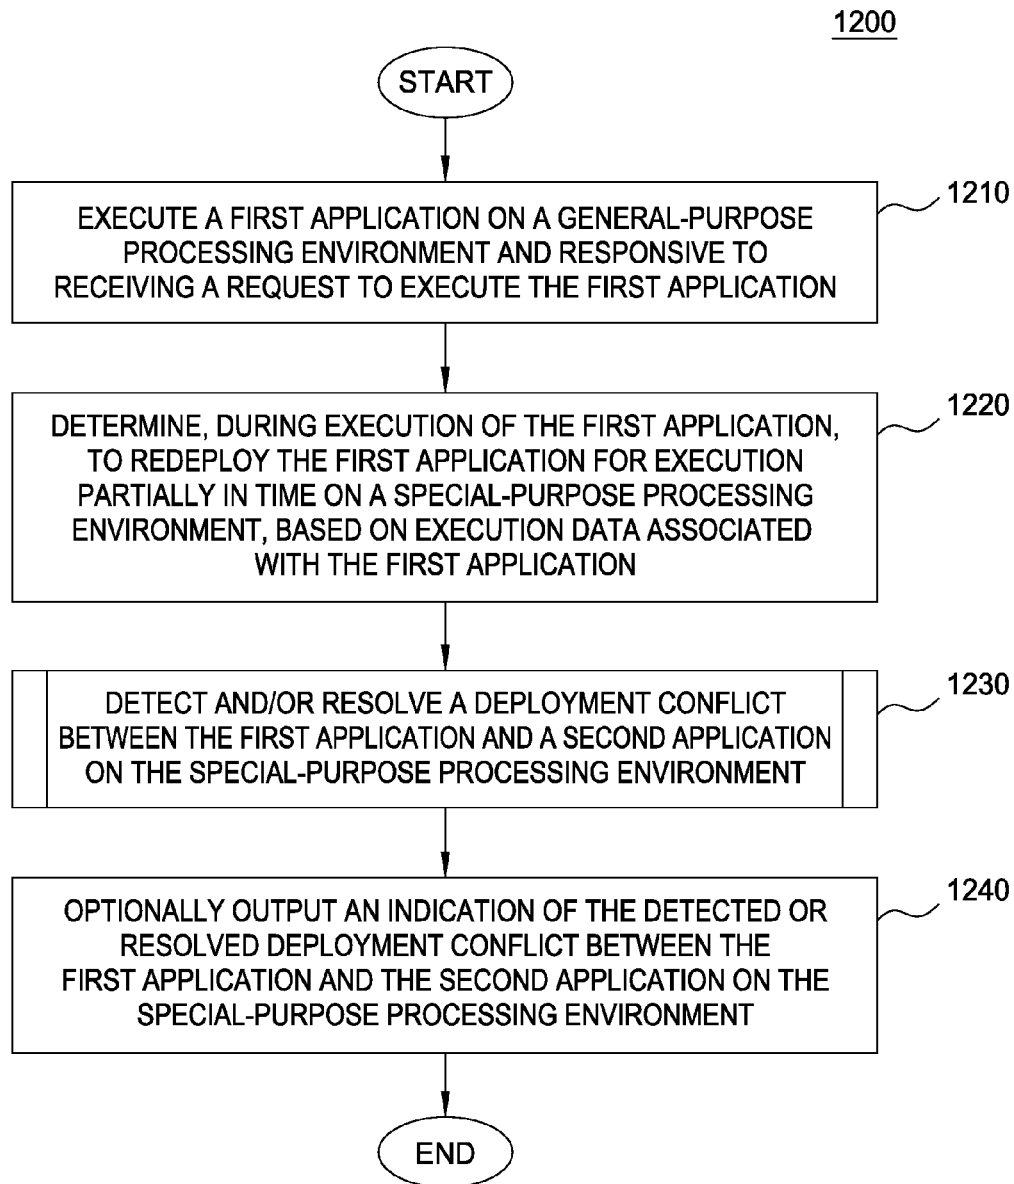


FIG. 12

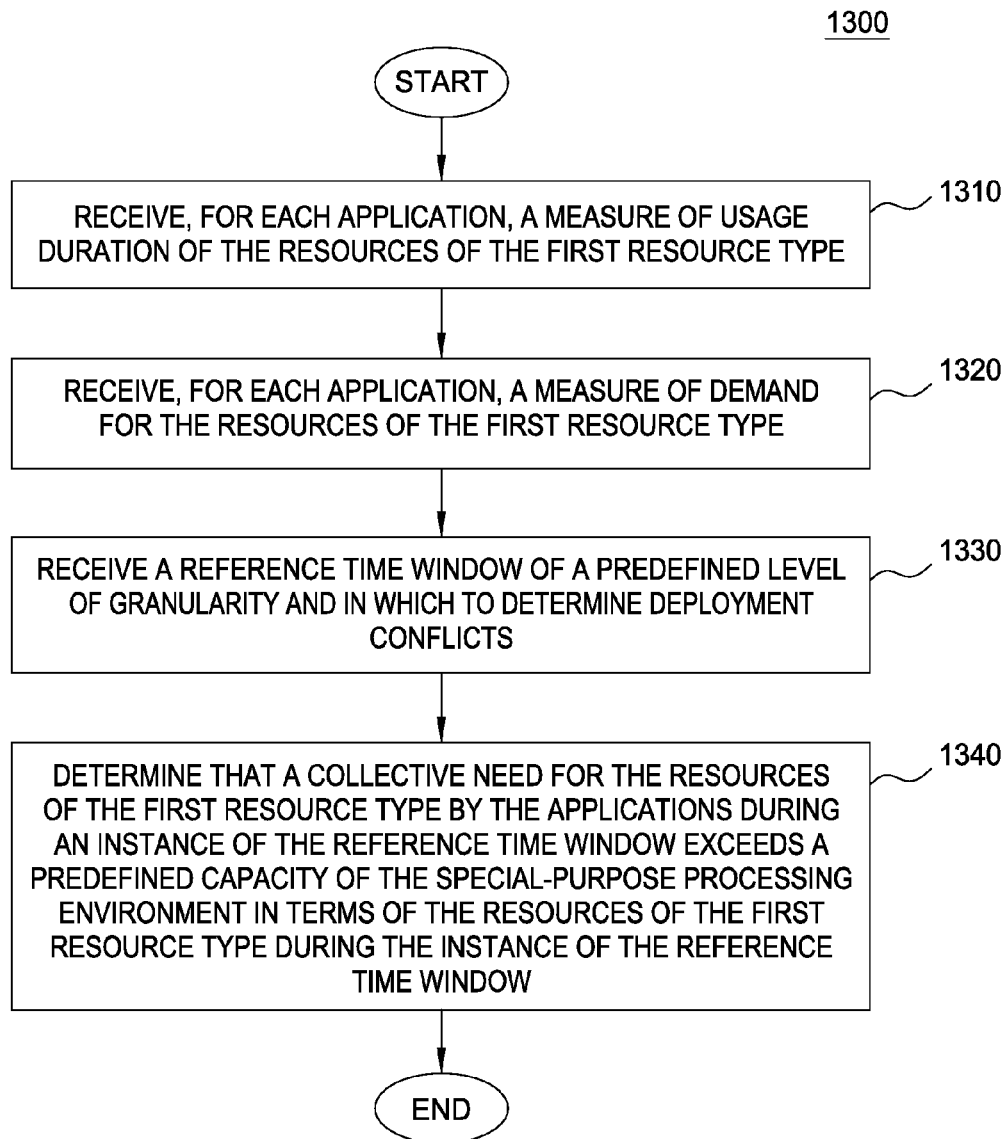


FIG. 13

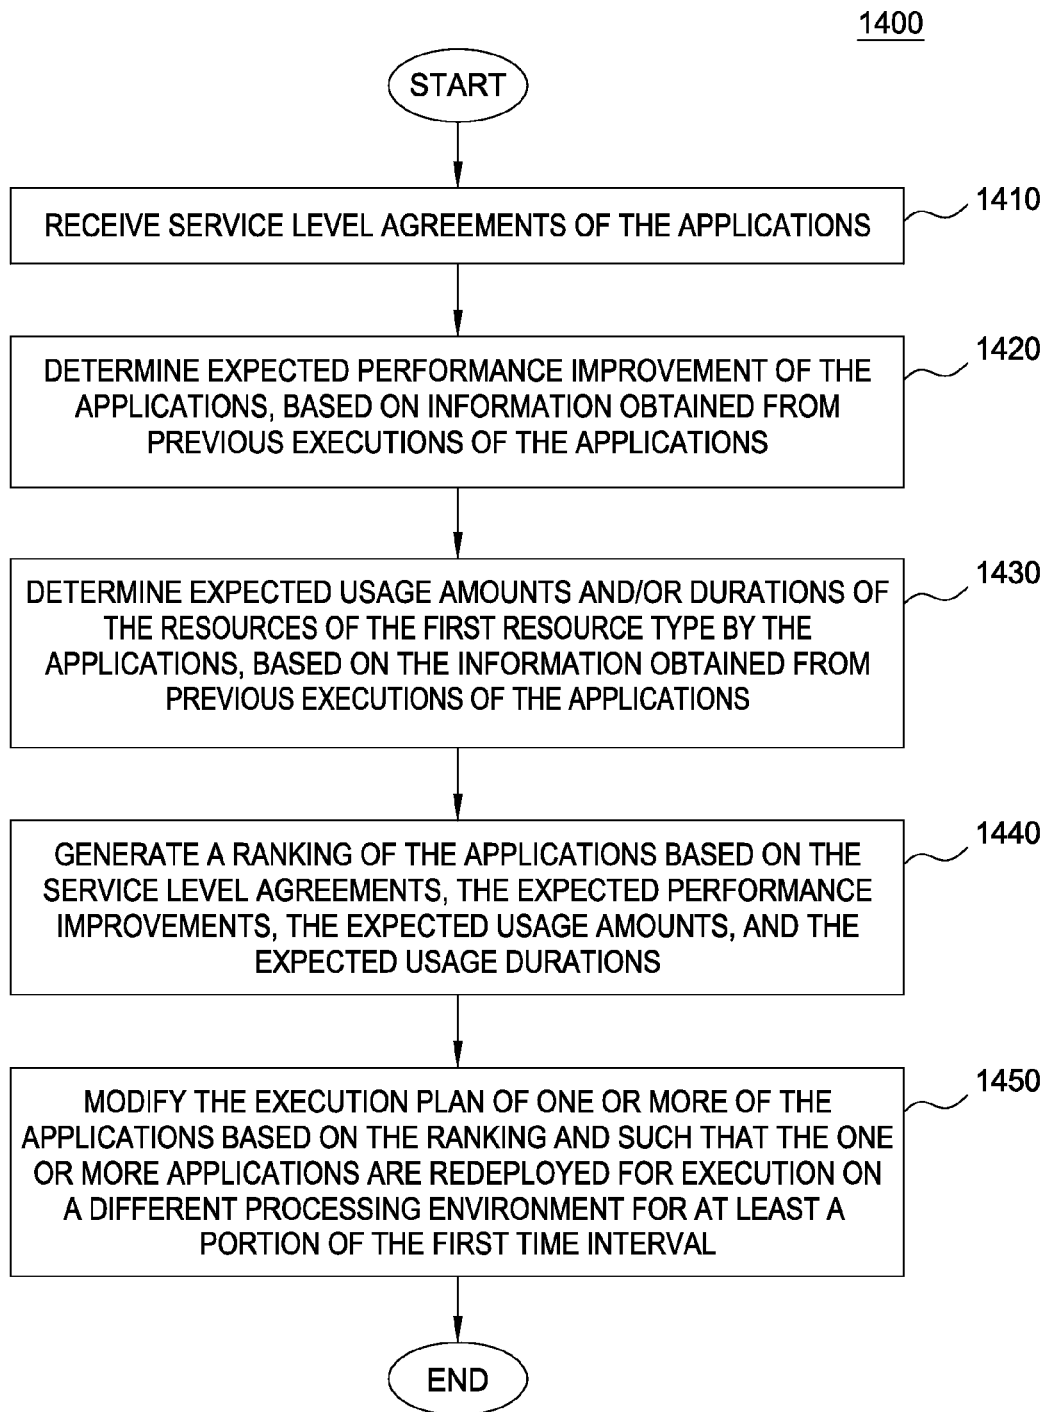


FIG. 14

RESOLVING DEPLOYMENT CONFLICTS IN HETEROGENEOUS ENVIRONMENTS

BACKGROUND

1. Field

Embodiments disclosed herein relate to application deployment. More specifically, embodiments disclosed herein relate to resolving deployment conflicts in heterogeneous environments.

2. Description of the Related Art

Advances in a virtualization technology and mainframe hardware have made a consolidation of applications onto a fewer number of centralized servers very attractive. However, a process of planning and performing such a migration is time-consuming, costly, and prone to an error. A risk is also involved in the migration process due to a high degree of complexity of applications deployed on server devices. For example, consider a large number of distributed applications that can be deployed on servers, also referred to as server devices. One server device could host an application on a middleware product to meet performance and high availability requirements. That application may be distributed over several server devices. This distributed application may require a database which is deployed on another server device. Messaging components to support information exchange with other internal or external applications may require a set of server devices and be deployed to several application servers. Effective consolidation of this set of servers may require an in-depth knowledge of configurations of the set of servers and an insight into how to best plan and run a migration of source applications to the final target platform.

SUMMARY

Embodiments presented in this disclosure provide a computer-implemented method that includes executing a first application in a first processing environment, responsive to receiving a request to execute the first application. The method also includes determining, during execution of the first application, to redeploy the first application for execution partially in time in a second processing environment, based on execution data associated with the first application. The second processing environment provides a higher capability than the first processing environment in terms of at least a first resource type. The method also includes resolving a deployment conflict between the first application and a second application in the second processing environment.

Other embodiments presented in this disclosure provide a computer program product that includes a computer-readable storage medium having program code embodied therewith. The program code is executable by one or more computer processors to execute a first application in a first processing environment, responsive to receiving a request to execute the first application. The program code is also executable to determine, during execution of the first application, to redeploy the first application for execution partially in time in a second processing environment, based on execution data associated with the first application. The second processing environment provides a higher capability than the first processing environment in terms of at least a first resource type. The program code is also executable to resolve a deployment conflict between the first application and a second application in the second processing environment.

Still other embodiments presented in this disclosure provide a system that includes one or more computer processors and a memory containing a program which, when executed

by the one or more computer processors, is configured to perform an operation that includes executing a first application in a first processing environment, responsive to receiving a request to execute the first application. The operation also includes determining, during execution of the first application, to redeploy the first application for execution partially in time in a second processing environment, based on execution data associated with the first application. The second processing environment provides a higher capability than the first processing environment in terms of at least a first resource type. The operation also includes resolving a deployment conflict between the first application and a second application in the second processing environment.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

So that the manner in which the above recited aspects are attained and can be understood in detail, a more particular description of embodiments of the invention, briefly summarized above, may be had by reference to the appended drawings.

It is to be noted, however, that the appended drawings illustrate only typical embodiments of this invention and are therefore not to be considered limiting of its scope, for the invention may admit to other equally effective embodiments.

FIG. 1 is a block diagram illustrating a cloud computing node, according to one embodiment presented in this disclosure.

FIG. 2 illustrates a cloud computing environment, according to one embodiment presented in this disclosure.

FIG. 3 illustrates abstraction model layers, according to one embodiment presented in this disclosure.

FIG. 4 is a block diagram of a networked system to evaluate an application to determine a processing environment suited for executing the application, according to one embodiment presented in this disclosure.

FIG. 5 is a block diagram showing heterogeneous processing environments on which the application may be deployed for execution, according to one embodiment presented in this disclosure.

FIG. 6 is a flowchart depicting a method for evaluating an application for deployment, according to one embodiment presented in this disclosure.

FIG. 7 is a flowchart depicting a method for using predefined conditions to evaluate an application for deployment, according to one embodiment presented in this disclosure.

FIG. 8 is a flowchart depicting a method for deploying an application for execution on the general-purpose processing environment, according to one embodiment presented in this disclosure.

FIG. 9 is a flowchart depicting a method for deploying an application for execution on a candidate special-purpose processing environment, according to one embodiment presented in this disclosure.

FIG. 10 is a flowchart depicting a method for deploying an application for execution on a target special-purpose processing environment, according to one embodiment presented in this disclosure.

FIG. 11 is a table characterizing execution plans for three different applications, according to one embodiment presented in this disclosure.

FIG. 12 is a flowchart depicting a method to manage deployment conflicts, according to one embodiment presented in this disclosure.

3

FIG. 13 is a flowchart depicting a method to identify deployment conflicts, according to one embodiment presented in this disclosure.

FIG. 14 is a flowchart depicting a method to resolve deployment conflicts, according to one embodiment presented in this disclosure.

DETAILED DESCRIPTION

Embodiments presented herein generally provide techniques to evaluate an application to determine a processing environment for executing the application and to detect and resolve deployment conflicts between applications. One embodiment provides a deployment tool configured to receive a request to execute the application. The deployment tool is further configured to deploy the application for execution on a general-purpose processing environment upon determining that insufficient execution data is present for the application. Otherwise (if sufficient execution data is present for the application), the deployment tool is further configured to deploy the application for execution partially in time on a special-purpose processing environment, based on the execution data. In one embodiment, executing an application partially in time on the special-purpose processing environment refers to the application executing for part or all of a duration of execution of the application (e.g., in cases where the application is being first deployed for execution) or part or all of a remaining duration of execution of the application (e.g., in cases where the application is being redeployed at runtime). If the application executes on the special-purpose processing environment for only part of a given duration, the application may execute on one or more other processing environments for the remainder of the given duration of execution of the application. The one or more other processing environments may include the general-purpose processing environment, other special-purpose processing environments, etc.

For example, the application may be a video streaming application that spends its first thirty seconds of execution primarily performing I/O operations to load a video from disk. Assume that the application spends the next sixty seconds of execution primarily performing network operations, e.g., transmitting a portion of the video to a client to allow client buffering. Further, the application spends the remainder of execution primarily performing graphics operations. To determine these resource demands of the application, the deployment tool analyzes execution data generated from executing the application on the general-purpose processing environment in order to generate an execution plan for the application. The execution plan may include one or more deployment rules, also referred to as migration rules, specifying where and/or when to deploy the application for execution.

Based on the execution plan, the deployment tool may deploy the application for thirty seconds of execution on a first special-purpose processing environment having enhanced I/O capabilities. The deployment tool may then migrate or redeploy the application during runtime to a second special-purpose processing environment having enhanced network capabilities, for a specified period of processing, such as sixty seconds of execution. To this end, the execution of the application on the first special-purpose processing environment may be suspended, and an associated runtime state of the application on the first special-purpose processing environment may be stored and transmitted to the second special-purpose processing environment, such that execution of the application on the second special-purpose processing environment may resume, on the second special-

4

purpose processing environment, from where execution of the application was previously suspended. The deployment tool may then redeploy the application during runtime for execution on a third special-purpose processing environment having enhanced graphics capabilities. At least in some embodiments, the execution plan may be generated based further on an execution plan template supplied by a user providing the application, the execution plan template specifying estimated resource demands of the application.

In some embodiments, part or all of the execution of the application may occur on the general-purpose processing environment rather than on a special-purpose processing environment, based on user preferences, which may be specified in a service level agreement associated with the application. Because a user providing the application may incur a higher monetary cost from executing the application on a special-purpose processing environment than on a general-purpose processing environment, the user may manage the monetary cost incurred via the service level agreement. Further, in some embodiments, the deployment tool may deploy the application for execution on a variety of special-purpose processing environments to generate trial data for the application. Such execution may also be referred to herein as trial execution. Depending on the embodiment, trial executions of the application on a given special-purpose processing environment may be from start to finish or only partial in time. The special-purpose processing environment may then be selected based on the trial data in addition to the execution data. For example, assume availability of two special-purpose processing environments with enhanced graphics capabilities, each environment having a graphics card from a different vendor. In one embodiment, the deployment tool may analyze the trial data to determine graphics card provides better performance for executing the application. Accordingly, when redeploying the application for execution on a special-purpose processing environment having enhanced graphics capabilities, the deployment tool may select the special-purpose processing environment having the determined graphics card.

Further, in some embodiments, a deployment conflict may be detected between applications and based on application-specific measures such as demand for resources of a given type, usage duration of resources of the given type, and a reference time window in which to determine conflicts. The reference time window may be a predefined time period or duration. At least in some embodiments, the deployment conflict may be resolved by modifying one or more execution plans of applications involved in the conflict, in order to redeploy at least one application to a different processing environment for at least part of the time window in which the deployment conflict is determined to occur. The deployment conflict may be resolved in favor of certain applications based on criteria such as service level agreements of the applications, expected performance improvement of the applications, and expected resource usage duration of the applications. Once the deployment conflict is detected and resolved, the applications may execute more efficiently in the processing environments at least in some cases.

In the following, reference is made to embodiments of the invention. However, it should be understood that the invention is not limited to specific described embodiments. Instead, any combination of the following features and elements, whether related to different embodiments or not, is contemplated to implement and practice the invention. Furthermore, although embodiments of the invention may achieve advantages over other possible solutions and/or over the prior art, whether or not a particular advantage is achieved by a given

5

embodiment is not limiting of the invention. Thus, the following aspects, features, embodiments and advantages are merely illustrative and are not considered elements or limitations of the appended claims except where explicitly recited in a claim(s). Likewise, reference to “the invention” shall not be construed as a generalization of any inventive subject matter disclosed herein and shall not be considered to be an element or limitation of the appended claims except where explicitly recited in a claim(s).

As will be appreciated by one skilled in the art, aspects of the present invention may be embodied as a system, method or computer program product. Accordingly, aspects of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a “circuit,” “module” or “system.” Furthermore, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this document, a computer readable storage medium may be any tangible medium that can contain, or store a program for use by or in connection with an instruction execution system, apparatus or device.

A computer readable signal medium may include a propagated data signal with computer readable program code embodied therein, for example, in baseband or as part of a carrier wave. Such a propagated signal may take any of a variety of forms, including, but not limited to, electro-magnetic, optical, or any suitable combination thereof. A computer readable signal medium may be any computer readable medium that is not a computer readable storage medium and that can communicate, propagate, or transport a program for use by or in connection with an instruction execution system, apparatus, or device.

Program code embodied on a computer readable medium may be transmitted using any appropriate medium, including but not limited to wireless, wireline, optical fiber cable, RF, etc., or any suitable combination of the foregoing.

Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like and conventional procedural programming languages, such as the “C” programming language or similar programming languages. The program code may execute entirely on the user’s computer, partly on the user’s computer, as a stand-alone software package, partly on the user’s computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the

6

remote computer may be connected to the user’s computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

Aspects of the present invention are described below with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

These computer program instructions may also be stored in a computer readable medium that can direct a computer, other programmable data processing apparatus, or other devices to function in a particular manner, such that the instructions stored in the computer readable medium produce an article of manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks.

The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which execute on the computer or other programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagrams in the Figures illustrate the architecture, functionality and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment or portion of code, which comprises one or more executable instructions for implementing the specified logical function(s). It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

It is understood in advance that although this disclosure includes a detailed description on cloud computing, implementation of the teachings recited herein are not limited to a cloud computing environment. Rather, embodiments of the present invention are capable of being implemented in conjunction with any other type of computing environment now known or later developed.

For convenience, the Detailed Description includes the following definitions which have been derived from the

“Draft NIST Working Definition of Cloud Computing” by Peter Mell and Tim Grance, dated Oct. 7, 2009.

Cloud computing is a model of service delivery for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, network bandwidth, servers, processing, memory, storage, applications, virtual machines, and services) that can be rapidly provisioned and released with minimal management effort or interaction with a provider of the service. This cloud model may include at least five characteristics, at least three service models and at least four deployment models.

Characteristics are as follows:

On-demand self-service: a cloud consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with the service’s provider.

Broad network access: capabilities are available over a network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops and PDAs).

Resource pooling: the provider’s computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to demand. There is a sense of location independence in that the consumer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state or datacenter).

Rapid elasticity: capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

Measured service: cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth and active user accounts). Resource usage can be monitored, controlled and reported providing transparency for both the provider and consumer of the utilized service.

Service Models are as follows:

Software as a Service (SaaS): the capability provided to the consumer is to use the provider’s applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based e-mail). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

Platform as a Service (PaaS): the capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including networks, servers, operating systems or storage, but has control over the deployed applications and possibly application hosting environment configurations.

Infrastructure as a Service (IaaS): the capability provided to the consumer is to provision processing, storage, networks and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems storage, deployed

applications, and possibly limited control of select networking components (e.g., host firewalls).

Deployment Models are as follows:

Private cloud: the cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on-premises or off-premises.

Community cloud: the cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy and compliance considerations). It may be managed by the organizations or a third party and may exist on-premises or off-premises.

Public cloud: the cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

Hybrid cloud: the cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

A cloud computing environment is service oriented with a focus on statelessness, low coupling, modularity and semantic interoperability. At the heart of cloud computing is an infrastructure comprising a network of interconnected nodes.

Referring now to FIG. 1, a schematic of an example of a cloud computing node is shown. Cloud computing node 10 is only one example of a suitable cloud computing node and is not intended to suggest any limitation as to the scope of use or functionality of embodiments of the invention described herein. Regardless, cloud computing node 10 is capable of being implemented and/or performing any of the functionality set forth hereinabove.

In cloud computing node 10 there is a computer system/server 12, which is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well-known computing systems, environments, and/or configurations that may be suitable for use with computer system/server 12 include, but are not limited to, personal computer systems, server computer systems, thin clients, thick clients, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputer systems, mainframe computer systems, and distributed cloud computing environments that include any of the above systems or devices, and the like.

Computer system/server 12 may be described in the general context of computer system-executable instructions, such as program modules, being executed by a computer system. Generally, program modules may include routines, programs, objects, components, logic, data structures, and so on that perform particular tasks or implement particular abstract data types. Computer system/server 12 may be practiced in distributed cloud computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed cloud computing environment, program modules may be located in both local and remote computer system storage media including memory storage devices.

As shown in FIG. 1, computer system/server 12 in cloud computing node 10 is shown in the form of a general-purpose computing device. The components of computer system/server 12 may include, but are not limited to, one or more processors or processing units 16, a system memory 28, and a bus 18 that couples various system components including system memory 28 to processor 16.

Bus 18 represents one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus and Peripheral Component Interconnects (PCI) bus.

Computer system/server 12 typically includes a variety of computer system readable media. Such media may be any available media that is accessible by computer system/server 12, and it includes both volatile and non-volatile media, removable and non-removable media.

System memory 28 can include computer system readable media in the form of volatile memory, such as random access memory (RAM) 30 and/or cache memory 32. Computer system/server 12 may further include other removable/non-removable, volatile/non-volatile computer system storage media. By way of example only, storage system 34 can be provided for reading from and writing to a non-removable, non-volatile magnetic media (not shown and typically called a "hard drive"). Although not shown, a magnetic disk drive for reading from and writing to a removable, non-volatile magnetic disk (e.g., a "floppy disk"), and an optical disk drive for reading from or writing to a removable, non-volatile optical disk such as a CD-ROM, DVD-ROM or other optical media can be provided. In such instances, each can be connected to bus 18 by one or more data media interfaces. As will be further depicted and described below, memory 28 may include at least one program product having a set (e.g., at least one) of program modules that are configured to carry out the functions of embodiments of the invention.

Program/utility 40, having a set (at least one) of program modules 42, may be stored in memory 28 by way of example, and not limitation, as well as an operating system, one or more application programs, other program modules, and program data. Each of the operating system, one or more application programs, other program modules and program data or some combination thereof, may include an implementation of a networking environment. Program modules 42 generally carry out the functions and/or methodologies of embodiments of the invention as described herein.

Computer system/server 12 may also communicate with one or more external devices 14 such as a keyboard, a pointing device, a display 24, etc.; one or more devices that enable a user to interact with computer system/server 12; and/or any devices (e.g., network card, modem, etc.) that enable computer system/server 12 to communicate with one or more other computing devices. Such communication can occur via I/O interfaces 22. Still yet, computer system/server 12 can communicate with one or more networks such as a local area network (LAN), a general wide area network (WAN), and/or a public network (e.g., the Internet) via network adapter 20. As depicted, network adapter 20 communicates with the other components of computer system/server 12 via bus 18. It should be understood that although not shown, other hardware and/or software components could be used in conjunction with computer system/server 12. Examples, include, but are not limited to: microcode, device drivers, redundant processing units, external disk drive arrays, RAID systems, tape drives, and data archival storage systems, etc.

Referring now to FIG. 2, illustrative cloud computing environment 50 is depicted. As shown, cloud computing environment 50 comprises one or more cloud computing nodes 10 with which local computing devices used by cloud consumers, such as, for example, personal digital assistant (PDA) or

cellular telephone 54A, desktop computer 54B, laptop computer 54C and/or automobile computer system 54N may communicate. Nodes 10 may communicate with one another. They may be grouped (not shown) physically or virtually, in one or more networks, such as Private, Community, Public, or Hybrid clouds as described hereinabove, or a combination thereof. This allows cloud computing environment 50 to offer infrastructure, platforms and/or software as services for which a cloud consumer does not need to maintain resources on a local computing device. It is understood that the types of computing devices 54A-N shown in FIG. 2 are intended to be illustrative only and that computing nodes 10 and cloud computing environment 50 can communicate with any type of computerized device over any type of network and/or network addressable connection (e.g., using a web browser).

Referring now to FIG. 3, a set of functional abstraction layers provided by cloud computing environment 50 (FIG. 2) is shown. It should be understood in advance that the components, layers and functions shown in FIG. 3 are intended to be illustrative only and embodiments of the invention are not limited thereto. As depicted, the following layers and corresponding functions are provided:

Hardware and software layer 60 includes hardware and software components. Examples of hardware components include mainframes, in one example IBM® zSeries® systems; RISC (Reduced Instruction Set Computer) architecture based servers, in one example IBM pSeries® systems; IBM xSeries® systems; IBM BladeCenter® systems; storage devices; networks and networking components. Examples of software components include network application server software, in one example IBM WebSphere® application server software; and database software, in one example IBM DB2C® database software. (IBM, zSeries, pSeries, xSeries, BladeCenter, WebSphere, and DB2 are trademarks of International Business Machines Corporation registered in many jurisdictions worldwide).

Virtualization layer 62 provides an abstraction layer from which the following examples of virtual entities may be provided: virtual servers; virtual storage; virtual networks, including virtual private networks; virtual applications and operating systems; and virtual clients.

In one example, management layer 64 may provide the functions described below. Resource provisioning provides dynamic procurement of computing resources and other resources that are utilized to perform tasks within the cloud computing environment. Metering and Pricing provide cost tracking as resources are utilized within the cloud computing environment, and billing or invoicing for consumption of these resources. In one example, these resources may comprise application software licenses. Security provides identity verification for cloud consumers and tasks, as well as protection for data and other resources. User portal provides access to the cloud computing environment for consumers and system administrators. Service level management provides cloud computing resource allocation and management such that required service levels are met. Service Level Agreement (SLA) planning and fulfillment provide pre-arrangement for, and procurement of, cloud computing resources for which a future requirement is anticipated in accordance with an SLA. The SLA generally specifies the services, priorities, responsibilities, guarantees and/or warranties that exist between a service provider and a customer.

Workloads layer 66 provides examples of functionality for which the cloud computing environment may be utilized. Examples of workloads and functions which may be provided from this layer include: mapping and navigation; software development and lifecycle management; virtual classroom

education delivery; data analytics processing; transaction processing; and mobile desktop.

Embodiments presented herein may be provided to end users through a cloud computing infrastructure. Cloud computing generally refers to the provision of scalable computing resources as a service over a network. More formally, cloud computing may be defined as a computing capability that provides an abstraction between the computing resource and its underlying technical architecture (e.g., servers, storage, networks), enabling convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction. Thus, cloud computing allows a user to access virtual computing resources (e.g., storage, data, applications, and even complete virtualized computing systems) in “the cloud,” without regard for the underlying physical systems (or locations of those systems) used to provide the computing resources.

Typically, cloud computing resources are provided to a user on a pay-per-use basis, where users are charged only for the computing resources actually used (e.g., an amount of storage space consumed by a user or a number of virtualized systems instantiated by the user). A user can access any of the resources that reside in the cloud at any time, and from anywhere across the Internet. In context of embodiments presented herein, a user may access deployed applications or related data available in the cloud. The deployed applications could execute on a suitable processing environment in the cloud, and a deployment tool could compute a cost to be charged for executing a given application, where the cost is determined based at least in part on the processing environment(s) on which the given application is deployed for execution. Further, deployment conflicts may be identified and resolved in the cloud. Thus, the user may access the applications from any computing system attached to a network connected to the cloud (e.g., the Internet) and be charged based on the processing environment(s) used.

FIG. 4 is a block diagram of a networked system 400 for evaluating an application for deployment, according to one embodiment presented in this disclosure. The networked system 400 includes a computer 402. The computer 402 may also be connected to other computers via the network 430. In general, the network 430 may be a telecommunications network and/or a wide area network (WAN). In a particular embodiment, the network 430 is the Internet.

The computer 402 generally includes a processor 404 connected via a bus 412 to a memory 406, a network interface device 410, a storage 408, an input device 414, and an output device 416. The computer 402 is generally under the control of an operating system. Examples of operating systems include UNIX, versions of the Microsoft Windows® operating system, and distributions of the Linux® operating system. More generally, any operating system supporting the functions disclosed herein may be used. The processor 404 is included to be representative of a single CPU, multiple CPUs, a single CPU having multiple processing cores, and the like. Similarly, the memory 406 may be a random access memory. While the memory 406 is shown as a single identity, it should be understood that the memory 406 may comprise a plurality of modules, and that the memory 406 may exist at multiple levels, from high speed registers and caches to lower speed but larger DRAM chips. The network interface device 410 may be any type of network communications device allowing the computer 402 to communicate with other computers via the network 430.

The storage 408 may be a persistent storage device. Although the storage 408 is shown as a single unit, the storage

408 may be a combination of fixed and/or removable storage devices, such as fixed disc drives, solid state disc (SSD) drives, floppy disc drives, tape drives, removable memory cards or optical storage. The memory 406 and the storage 408 may be part of one virtual address space spanning multiple primary and secondary storage devices.

The input device 414 may be any device for providing input to the computer 402. For example, a keyboard or mouse may be used. The output device 416 may be any device for providing output to a user of the computer 402. For example, the output device 416 may be any conventional display screen or set of speakers, along with their respective interface cards, i.e., video cards and sound cards (not shown). Although shown separately from the input device 414, the output device 416 and input device 414 may be combined. For example, a display screen with an integrated touch-screen may be used.

As shown, the memory 406 of the computer 402 includes a deployment tool 450. The storage 408 of the computer 402 includes an application registry 452, a service level agreement 454, an execution plan template 456, execution data 458, trial data 460, and an execution plan 462. In one embodiment, the application registry 452 specifies a set of applications available for deployment by the deployment tool 450. The deployment tool 450 may deploy an application for execution on a general-purpose processing environment to generate the execution data 458. At least in some embodiments, the general-purpose processing environment may be selected from multiple general-purpose processing environments having similar capabilities in terms of resource types. The execution data 458 characterizes runtime workloads or resource demands of the application over time and in terms of at least one resource type, when executing on the general-purpose processing environment.

In one embodiment, the deployment tool 450 deploys the application for execution partially in time on a special-purpose processing environment selected based on the execution data 458. For example, the application may execute on the special-purpose processing environment for a first time duration and on the general-purpose processing environment (or another special-purpose processing environment) for a second time duration. Each time duration may be also referred to herein as a time slice or time interval. The selected special-purpose processing environment may be configured to better suit the resource demands of the application as indicated by the execution data, compared to the general-purpose processing environment. In some embodiments where distinct special-purpose processing environments—referred to as candidate environments—suit the resource demands of the application, the deployment tool 450 may deploy the application to execute on each of the candidate environments to generate trial data 460. The trial data 460 characterizes runtime performance of the application over time and when executing in a special-purpose processing environment. The deployment tool 450 may then deploy the application for execution partially in time on one of the candidate environments—referred to as a target environment—that is selected based on the trial data 460 as best suiting the resource demands of the application, as compared to the other candidate environments. In some embodiments, the deployment tool 450 generates the trial data 460 and/or the execution data based at least in part on output obtained from predefined application tuning and/or benchmarking utilities (e.g., top, perfmonit, nmon, etc).

At least in some embodiments, the application stores data specifying where and/or when to deploy the application in the form of the execution plan 462 for the application. The execution plan 462 may be generated and/or updated based on the

13

execution data **458** and/or the trial data **460** and optionally based further on an execution plan template **456** supplied by a user providing the application, where the execution plan template **456** specifies expected resource demands of the application. The execution plan template **456** may be supplied to assist the deployment tool **450** to select a suitable processing environment more accurately and/or quickly. Further, an execution plan that is generated based on the execution data **458** and not the trial data **460** may be referred to herein as a baseline execution plan. At least in some embodiments, the target environment and/or the candidate environments and/or deployment intervals may be selected based further on user preferences pertaining to processing environments and/or resource types and/or monetary costs thereof. At least in some embodiments, the user preferences may be specified in the service level agreement **454**. Put another way, a user providing the application may specify, in the service level agreement, how much the user is willing to pay for executing the application, in terms of monetary costs. In still other words, a user may specify a quality of service desired by the user for executing the application.

In some embodiments, in addition to quantifying the willingness of the user to pay in terms of monetary cost, the service level agreement may qualify the willingness of the user to pay in terms of time of day, day of week, enhanced capability provided, etc. For example, the user may express, in the service level agreement, a willingness to pay extra for enhanced graphics capabilities, but only on weekends between 5-8 pm. Additionally or alternatively, the service level agreement may also specify a willingness to pay extra for enhanced network capabilities on Mondays through Fridays between 5-8 pm. The service level agreement may also specify time periods and/or enhanced capabilities for which the user is not willing to pay extra. The deployment tool **450** may deploy the application for execution on the general-purpose processing environment during such time periods. Further, in some embodiments, the execution plan may also be modified based on the service level agreement **454**.

Accordingly, the application may execute more efficiently using—even if only partially in time—the enhanced capabilities provided by one or more special-purpose processing environments. Further, the user providing the application may adjust the service level agreement to limit or otherwise manage the monetary cost incurred from executing the application. Consequently, the user need not endure performance associated with executing the application from start to finish on a general-purpose processing environment, which may be undesirably slow to the user. Further, the user need not incur a monetary cost associated with executing the application from start to finish on a special-purpose processing environment, which may be prohibitively expensive for the user. The user also need not incur an effort involved in attempting to determine, with some degree of accuracy, special-purpose processing environment(s) suited to the resource demands of the application. Further, processing environments with distinct capabilities in a heterogeneous computing environment may also be utilized more effectively, because applications are migrated to those processing environments most equipped to satisfy the resource demands of the application—as the resource demands change over time during execution of the application.

FIG. 5 is a block diagram showing heterogeneous processing environments on which an application may be deployed for execution, according to one embodiment presented in this disclosure. As shown, the deployment tool **450** deploys applications **502** to processing environments **504** based on execu-

14

tion plans **462**, service level agreements **454**, and/or execution plan templates **456**. The processing environments **504** include a general-purpose processing environment **504₃**, a special-purpose processing environment **504₁** having enhanced storage capability, a special-purpose processing environment **504₂** having enhanced graphics capability, a special-purpose processing environment **504₄** having enhanced audio capability, and a special-purpose processing environment **504₅** having enhanced network capability. The processing environments **504** may also include other environments **504₆**, such as special-purpose processing environments with enhanced processor or memory capability, etc.

In one embodiment, the deployment tool **450** performs a selection operation **506** to determine when and/or where to deploy (and redeploy) the applications **502** to the processing environments **504**, on the basis of the execution plans **462**, the service level agreements **454**, and/or the execution plan templates **456**. As described above, to this end, the execution plan may include one or more deployment rules specifying time durations and target processing environments, during and on which the application is to execute. The time durations may be measured relative to a time at which execution of (a given instance of) the application is commenced. In some embodiments, the service level agreement may also specify part or all of the information represented by the one or more redeployment rules or may include the one or more redeployment rules outright.

In one embodiment, the deployment tool **450** also performs an analysis operation **508** on execution data and/or trial data generated from executing the applications **502**, in order to generate and/or update the execution plans **462**. Accordingly, the application may be deployed in a manner that better suits the needs of the application and/or the user providing the application, relative to alternative approaches for managing execution of the application on heterogeneous processing environments.

In some embodiments, the deployment tool **450** may also perform one or more predefined operations to reduce overhead incurred from redeploying the application from processing environment to processing environment. For example, assume that the deployment tool **450** is to redeploy an application from a current processing environment to a next processing environment. Assume further that there are two candidate processing environments having the same capability. In one embodiment, to reduce redeployment overhead, the deployment tool **450** may favor the candidate processing environment more proximate to the current processing environment in terms a set of predefined criteria such as network distance, geographical distance, etc. In cases where the deployment tool **450** determines that any processing benefit gained from redeploying the application is outweighed by the redeployment overhead incurred, the deployment tool **450** may refrain from one or more redeployments of the application. The deployment tool **450** may also update the execution plan accordingly, such as to reduce the number of redeployments specified in the execution plan. In other embodiments, a requesting entity and/or an administrative user of the deployment tool **450** may also specify a maximum number of redeployments for a given execution of the application. In some embodiments, the deployment tool **450** may also generate an alert to inform the administrative user, upon determining that performance of the application has degraded beyond a predetermined threshold, e.g., a maximum execution time for the application. Based on the alert, the administrative user may reconfigure the deployment tool **450**, the application, the execution plan, the network topology, etc.

15

FIG. 6 is a flowchart depicting a method 600 for evaluating an application for deployment, according to one embodiment presented in this disclosure. As shown, the method 600 begins at step 610, where the deployment tool 450 receives a request to execute the application. At step 620, upon determining that insufficient execution data is present for the application, the deployment tool 450 deploys the application for execution on a general-purpose processing environment. Depending on the embodiment, whether sufficient execution data is present may be determined based on a predefined condition, such as the application having been executed a threshold count of times on the general-purpose processing environment, the application having been executed for a threshold time duration on the general-purpose processing environment, a threshold amount of the execution data having been generated, etc. The thresholds may be user-specified as desired to suit the needs of a particular case. Depending on the embodiment, the predefined condition may specify qualitative and/or quantitative criteria pertaining to the execution data.

At step 630, upon determining that sufficient execution data and insufficient trial data exists for the application, the deployment tool 450 deploys the application for execution on one or more candidate environments, each selected from a set of available special-purpose processing environments. Depending on the embodiment, whether sufficient trial data is present for a given special-purpose processing environment may be determined based on a predefined condition, such as the application having been executed a threshold count of times, the application having been executed for a threshold time duration, a threshold amount of the trial data having been generated, etc. Depending on the embodiment, the predefined condition may specify qualitative and/or quantitative criteria pertaining to the trial data. Depending on the embodiment, the predefined condition may pertain to a given special-purpose processing environment, all candidate environments, all available special-purpose processing environments, etc.

At step 640, upon determining that sufficient execution data and sufficient trial data are present for the application, the deployment tool 450 deploys the application for execution at least partially in time on a target environment selected from the one or more candidate environments based on the trial data—e.g., for a time duration determined based on the execution data and/or the trial data. For example, the time duration may be determined based on corresponding time durations during which the trial data indicates that the application is particularly demanding in terms of a given resource type, e.g., one for which enhanced capability is provided by the target environment, relative to the general-purpose processing environment. In particular, the time duration may be determined by any statistical measure of the corresponding time durations, such as an arithmetic mean. After the step 640, the method 600 terminates.

In some embodiments, each execution of the application occurs only upon receipt of an explicit request to execute the application from a requesting entity such as a user of the deployment tool. In such embodiments, a suitable processing environment may not necessarily be determined until multiple requests for execution of the application are received and processed. This is because accuracy of the execution data and/or trial data may often only improve gradually as the execution data and/or trial data is updated based on each subsequent execution of the application. In other embodiments, one or more executions of the application occur independent of any such request from the requesting entity, to facilitate selection of a suitable processing environment when the request is subsequently received. In still other embodiments, upon receiving a request for a single execution of the

16

application, the deployment tool 450 nevertheless deploys the application for multiple executions, such that a suitable processing environment may sooner be determined.

FIG. 7 is a flowchart depicting a method 700 for using predefined conditions to evaluate an application for deployment, according to one embodiment presented in this disclosure. As shown, the method 700 begins at step 710, where the deployment tool 450 receives a request to execute the application. At step 720, the deployment tool 450 determines whether the application has been executed on the general-purpose processing environment for a threshold count of times. If not, the deployment tool 450 deploys the application for execution on the general-purpose processing environment (step 725). The step 725 is further described below in conjunction with FIG. 8.

On the other hand, if the deployment tool 450 determines that the threshold is met (step 720), then the deployment tool 450 determines whether trial executions are enabled (step 730). For example, whether trial executions are enabled may be specified by a flag set by a user providing the application and/or by an administrative user managing the deployment tool 450. If trial executions are enabled, then the deployment tool 450 determines whether the application has been executed on special-purpose processing environments for a threshold count of times (step 745). If not, the deployment tool 450 deploys the application for execution on a candidate special-purpose processing environment (step 750). The step 750 is further described below in conjunction with FIG. 10.

If trial executions are not enabled (step 730) or the application has not been executed on the special-purpose processing environments for the threshold count of times (step 745), then the deployment tool 450 determines whether an execution plan and/or a service level agreement is present for the application, with no prohibition on special-purpose processing (step 735). For example, in some embodiments, the service level agreement may specify that a user has prohibited special-purpose processing.

If the execution plan and/or the service level agreement is present for the application, with no prohibition on special-purpose processing, then the deployment tool 450 deploys the application for execution at least partially in time on a target special-purpose processing environment selected based on the execution plan and/or the service level agreement (step 760); otherwise, the deployment tool 450 deploys the application for execution on the general-purpose processing environment (the step 725). The step 760 is further described below in conjunction with FIG. 9. Depending on the embodiment, the target special-purpose processing environment may be selected from the candidate environments or directly from the set of available special-purpose processing environments. The duration for which the application executes on the special-purpose processing environment is also determined based on the execution plan and/or the service level agreement. After the steps 725, 750 or 760, the method 700 terminates.

FIG. 8 is a flowchart depicting a method 800 for deploying an application for execution on the general-purpose processing environment, according to one embodiment presented in this disclosure. The method 800 corresponds to the step 725 of FIG. 7. As shown, the method 800 begins at step 810, where a general-purpose hardware environment is selected by the deployment tool 450 and/or the requesting entity. At step 820, the deployment tool 450 analyzes execution data generated from executing the application on the general-purpose processing environment. At step 830, the deployment tool 450 increments a counter representing a total number of times the

17

application has executed, e.g., from start to finish, on the general-purpose processing environment.

At step **840**, the deployment tool **450** generates or updates an execution plan for the application, based on analyzing the execution data. At step **850**, the deployment tool **450** outputs the execution plan in whole or in part to the requesting entity. In alternative embodiments, the deployment tool **450** generates output for the requesting entity based on the execution plan, while the execution plan itself is not output to the requesting entity. For example, the deployment tool **450** may generate output in the form of suggested special-purpose processing environments and associated monetary costs, and the output may be generated based on the execution plan and/or service level agreement. The user may specify to adopt one or more of the suggested special-purpose processing environments for subsequent executions of the application. The execution plan, service level agreement, and/or execution plan template may be modified to reflect the user adoption. Accordingly, the output assists the user in discovering when and where to execute the application on a special-purpose processing environment, to improve a cost effectiveness of executing the application for the user at least in some cases. In alternative embodiments, the deployment tool **450** determines to adopt the execution plan based on information included in the service level agreement, e.g., monetary cost thresholds, without prompting for any input from the user. After the step **850**, the method **800** terminates.

FIG. 9 is a flowchart depicting a method **900** for deploying an application for execution on a candidate special-purpose processing environment, according to one embodiment presented in this disclosure. The method **900** corresponds to the step **750** of FIG. 7. As shown, the method **900** begins at step **910**, where the candidate special-purpose processing environment is selected by the deployment tool **450** and/or the requesting entity. At step **920**, the deployment tool **450** analyzes trial data generated from executing the application on the candidate special-purpose processing environment. At step **930**, the deployment tool **450** increments a counter representing a total number of times the application has executed, e.g., from start to finish, on the special-purpose processing environment. At step **940**, the deployment tool **450** updates the execution plan based on analyzing the trial data. If the execution plan specifies to perform any further execution of the application (step **950**), then the method **900** returns to the step **920** to further execute the application. Otherwise, the deployment tool **450** outputs, to the requesting entity, the execution plan itself or an output generated from the execution plan (step **960**). After the step **960**, the method **900** terminates.

FIG. 10 is a flowchart depicting a method **1000** for deploying an application for execution on a target special-purpose processing environment, according to one embodiment presented in this disclosure. The method **1000** corresponds to the step **760** of FIG. 7. As shown, the method **1000** begins at step **1010**, where the deployment tool **450** determines a target special-purpose processing environment based on the execution plan and/or service level agreement. At step **1020**, the deployment tool **450** analyzes data generated from executing the application on the target special-purpose processing environment. The generated data may be referred to herein as target data and, as with trial data, may also be used by the deployment tool **450** in modifying the execution plan (step **1030**). At step **1040**, the deployment tool **450** determines whether the execution plan specifies to perform any further execution of the application, and if so, the method **1000** returns to the step **1010** to further execute the application. Otherwise, the deployment tool **450** outputs, to the requesting

18

entity, the execution plan itself or an output generated from the execution plan (step **1050**). After the step **1050**, the method **1000** terminates.

Although embodiments are described herein with reference to analyzing and deploying applications based on execution data and/or trial data, other embodiments are broadly contemplated. For example, in one embodiment, an application may be analyzed and/or deployed based on further data characterizing resource needs of the application, where the data is generated by analyzing code of the application, also referred to as static analysis. The data is also referred to herein as static data. Examples of the code include source code, object code, and machine code. In contrast to execution data, which is generated from executing the application on a general-purpose processing environment, and trial data, which is generated from executing the application on a candidate special-purpose processing environment, static data is generated without executing the application on any processing environment. When analyzing code of the application, the deployment tool **450** also analyze any associated code, such as application programming interfaces (APIs), libraries, modules, etc.

In one embodiment, the static data may characterize resource demands of an application in terms of at least one resource type. For example, assume that upon analysis, the source code of the application includes function calls to a socket API. Accordingly, the deployment tool **450** may determine that the application demands network resources. The number and nature of the function calls may also be further analyzed, according to predefined rules specified by a user of the deployment tool **450**, to determine a likely extent to which network resources are demanded by the application. Further, code analysis may be performed with assistance of one or more predetermined code browsing tools, such as included in Cscope and some integrated development environment (IDE) tools such as Eclipse™. Depending on the embodiment, code analysis may be programmatically and/or manually performed in part or in full. In some embodiments, a report is programmatically generated based on code analysis, the report summarizing resource demands of the application. In such embodiments, the static data may further include the generated report.

Further, in one embodiment, the static data may serve as a substitute for part or all of execution data and/or trial data. For example, where the amount of execution data would otherwise be insufficient, the execution data may nevertheless suffice—when combined a sufficient amount of static data—to allow the application to be deployed for execution partially in time on a special-purpose processing environment and, to that end, generate a set of candidate special-purpose processing environments. As another example, where the amount of trial data would otherwise be insufficient, the trial data may nevertheless suffice—when combined with a sufficient amount of static data—to allow the application to be deployed for execution partially in time on a special-purpose processing environment selected from the set of candidate special-purpose processing environments. Further still, the special-purpose processing environment and/or the candidate environments may be selected based on the static data. The amount of relative weight accorded to each of the execution data, the trial data, and the static data—in terms of analysis, deployment, and/or data sufficiency determination—may be configured by a user of the deployment tool **450** and may be tailored to suit the needs of an individual case.

In one embodiment, the application is deployed based further on a current utilization of the special-purpose processing environment, where the current utilization characterizes a

workload of the special-purpose processing environment in terms of at least one resource type. For example assume that the deployment tool **450** is to redeploy the application to a special-purpose processing environment with enhanced network capabilities. Upon determining that the special-purpose processing environment has a current utilization exceeding a user-specified threshold, the deployment tool **450** refrains from redeploying the application to the special-purpose processing environment after all, to avoid over-utilization of the special-purpose processing environment. In some embodiments, executing the application **450** partially in time on a given special-purpose environment may incur a monetary cost that is dependent on a current utilization of the special-purpose environment. In some embodiments, the deployment tool **450** may nevertheless deploy the application for execution on the special-purpose processing environment, upon determining that the service level agreement associated with the application indicates that the user providing the application is willing to pay the higher monetary cost associated with a high level of utilization of the special-purpose processing environment.

Further, in some embodiments, deployment conflicts—also referred to herein as scheduling conflicts—may arise between applications executing and/or to be executed in the processing environments. Such deployment conflicts may be identified and/or resolved using the techniques disclosed herein. For example, to detect a deployment conflict between multiple execution plans, the deployment tool **450** may operate as follows. Although embodiments are described herein with reference to operations of the deployment tool **450**, some or all of the operations may be performed by a scheduler component of the deployment tool **450**, a standalone scheduling application, etc. In one embodiment, when starting an application for which an execution plan is already generated, the deployment tool **450** may perform a first step involving evaluating the plan to determine each specialized hardware component or resource type that the application uses during execution. For each specialized hardware component or resource type, the deployment tool **450** determines the overall capacity for the respective specialized hardware component or resource type. For instance, if the resource type is network resources, the deployment tool **450** may determine a total bandwidth of all network adapters available for use in the specialized processing environment.

In one embodiment, the deployment tool **450** may then iterate through all executing applications having execution plans and estimate, based on previous executions of those applications, when each execution is likely to use the specialized hardware component or resource type. For example, the estimations may be based on the execution data **458** and the trial data **460** of FIG. 1. In some embodiments, the above determinations may be repeated for additional processing environments and/or additional specialized hardware components or resource types. In a particular embodiment, only the specialized hardware components or resource types identified in the first step are considered. As an example, the estimations may be determined based on a planned start time for using the specialized hardware component or resource type, duration of the planned usage, and capacity required by the application. For example, if the resource type is network resources, then the capacity may refer to bandwidth required by the application. Depending on the embodiment, the capacity may be represented as a predefined statistical measure such as a maximum capacity, an average capacity, etc.

In one embodiment, the deployment tool **450** may then determine, for each specialized hardware component or resource type used by the application, a set of other executing

applications that use or are likely to use the same specialized hardware component or resource type during the same time window. The deployment tool **450** may then determine whether the collective needs of the applications exceed an overall capacity of the specialized hardware component or resource type. If so, then the deployment tool **450** has identified a deployment conflict among the applications. At least in some embodiments, the granularity of the time window may be tailored to suit the needs of a particular case. Further, where there are multiple special-purpose processing environments in the cloud, and individual applications are not readily balanced across such processing environments, then the deployment tool **450** may determine if the capacity may be reached for all possible combinations of the applications and processing environments. Such determinations may be made using predefined resource allocation or combinatorial optimization techniques, such as knapsack problem techniques.

Although the operations of the deployment tool **450** are described above as occurring when an application commences execution, in alternative embodiments, the operations may also occur when a predefined trigger condition is satisfied. In one embodiment, the trigger condition is satisfied when the application sufficiently deviates from its normal routine in terms of resource usage, as compared to previous executions of the application. Thus, the planned time windows in which the application is to use specialized hardware components or resource types may need to be rescheduled. In some embodiments, only specialized hardware needs of the application that occur subsequent to the current point in time of the deviating application are evaluated by the deployment tool **450**.

FIG. 11 is a table **1100** characterizing execution plans for three different applications, according to one embodiment presented in this disclosure. At least in some embodiments, deployment conflicts between applications may be more readily identified when the execution plans of the applications overlap in time in terms of resource use and resource type and when execution of the applications are commenced at or around the same time. As shown, the execution plans include a first execution plan **1102** for a media player application, a second execution plan **1104** for a video game, and a third execution plan **1106** for a shopping application.

Assume that the media player application, when executed, exhibits a pattern of ten seconds of I/O-bound behavior (e.g., buffering data) followed by sixty seconds of graphics-bound behavior (e.g., playing a video) and in turn followed by two seconds of I/O-bound behavior (e.g., displaying related videos). Assume further that the video game, when executed, exhibits a pattern of twenty seconds of I/O-bound behavior (e.g., loading game assets) followed by ten to sixty seconds of user input processing (e.g., user selection of options in game menus) and in turn followed by an arbitrarily lengthy amount of graphics-bound behavior (e.g., playing the game). Assume also that the shopping application, when executed, exhibits a pattern of one second of I/O-bound behavior (e.g., loading application assets) followed by thirty to ninety seconds of user input processing (e.g., user selection of desired products or services) and in turn followed by ten to twenty seconds of general purpose processing (e.g., processing an order placed by the user).

In some embodiments, when such applications contend for I/O, the deployment tool **450** may modify the third execution plan **1106** such as to use I/O resources from a general-purpose processing environment, at least in part because the duration required for the I/O operations is minimal in this example. Put another way, the performance improvements gained from the shopping application using I/O resources from a special-

21

purpose I/O processing environment would likely not justify the overhead in redeploying the shopping application to the special-purpose I/O processing environment.

In some embodiments, the deployment tool **450** may detect and resolve deployment conflicts between applications, regardless of whether the applications are executing or are to be executed. Further, the deployment conflicts may be detected and resolved regardless of whether the applications commence execution at the same time. For example, suppose that the media player application is not launched until the video game is already executing and consuming a greater amount of resources than when the video game first commenced execution.

At least in such scenarios, the deployment tool **450** may resolve deployment conflicts between the applications based on predefined criteria. Examples of the predefined criteria include measures such as available resources in a given processing environment, expected remaining resource usage amounts and durations of the executing application, and expected resource usage amounts and durations of the launching application. The predefined criteria may also include additional measures such as an expected remaining performance improvement of the executing application and an expected performance improvement of the launching application. The predefined criteria may further include respective overheads of redeploying the executing application and launching application, respectively. The predefined criteria may still further include service level agreements of the applications and differences therebetween, whether a predefined fairness algorithm, such as round-robin scheduling, is currently in use, etc. In some embodiments, the deployment tool **450** may also generate a ranking of the applications based on the predefined criteria. Depending on the embodiment, the ranking may be specific to a given processing environment or may not be specific to any single processing environment. The ranking may be determined based on resource allotment and/or execution timeline. The deployment tool **450** may then modify one or more execution plans based on the ranking, in order to redeploy one or more applications and resolve the deployment conflict. At least in some cases, if applications are deployed at difference points in time, allowing a previously deployed application to complete may be more efficient, from a resource utilization perspective, than redeploying the application prematurely to a different processing environment. By recognizing individual resource needs over the lifespans of the applications, the deployment tool **450** may generate modified execution plans that result in a more efficient deployment schedule at least in some cases.

FIG. **12** is a flowchart depicting a method **1200** to manage deployment conflicts, according to one embodiment presented in this disclosure. As shown, the method **1200** begins at step **1210**, where the deployment tool **450** executes a first application on a general-purpose processing environment and responsive to receiving a request to execute the first application. At step **1220**, the deployment tool **450** determines, during execution of the first application, to redeploy the first application for execution partially in time on a special-purpose processing environment, based on execution data associated with the first application. At step **1230**, the deployment tool **450** detects and/or resolves a deployment conflict between the first application and at least a second application on the special-purpose processing environment. The step **1230** is further described below in conjunction with the methods **1300**, **1400** of FIGS. **13-14**. At step **1240**, the deployment tool **450** optionally outputs an indication of the detected and/or resolved deployment conflict between the first application

22

and the second application on the special-purpose processing environment. After the step **1240**, the method **1200** terminates.

FIG. **13** is a flowchart depicting a method **1300** to identify deployment conflicts, according to one embodiment presented in this disclosure. The method **1300** corresponds to at least part of the step **1240** of FIG. **12**. As shown, the method **1300** begins at step **1310**, where the deployment tool **450** receives, for each application, a measure of usage duration of the resources of the first resource type. At step **1320**, the deployment tool **450** receives, for each application, a measure of demand for the resources of the first resource type. At step **1330**, the deployment tool **450** receives a reference time window of a predefined level of granularity and in which to determine deployment conflicts. At step **1340**, the deployment tool **450** determines that a collective need for the resources of the first resource type by the applications during an instance of the reference time window exceeds a predefined capacity of the special-purpose processing environment. The predefined capacity pertains to the resources of the first resource type during the instance of the reference time window. At least in some embodiments, the determination at the step **1340** corresponds to a deployment conflict between the applications during the instance of the reference time window. Accordingly, the deployment conflict is identified by the deployment tool **450**. After the step **1340**, the method **1300** terminates.

FIG. **14** is a flowchart depicting a method **1400** to resolve deployment conflicts, according to one embodiment presented in this disclosure. The method **1300** corresponds to at least part of the step **1240** of FIG. **12**. As shown, the method **1400** begins at step **1410**, where the deployment tool **450** receives service level agreements associated with the applications. At step **1420**, the deployment tool **450** determines expected performance improvements of the application, such as based on information obtained from previous executions of the applications, e.g., the execution data **458** and the trial data **460** of FIG. **1**. At step **1430**, the deployment tool **450** determines expected usage amounts and/or durations of the resources of the first resource type by the applications, based on the information obtained from previous executions of the applications. At least in some embodiments, determinations of the expected performance improvements, usage amounts, and usage durations may be based further on one or more predefined estimation rules.

At step **1440**, the deployment tool **1450** generates a ranking of the applications based on criteria including the service level agreements, the expected performance improvements, the expected usage amounts, and the expected usage durations. For example, the ranking may be generated based on a predefined function of the criteria and may include weighting associated with each criterion. At step **1450**, the deployment tool **1450** modifies the execution plan of one or more of the applications based on the ranking. In some embodiments, the execution plan is modified such that the one or more applications are redeployed for execution on a different processing environment for at least a portion of the first time interval. For instance, the different environment may be the general-purpose processing environment or another special-purpose processing environment. After the step **1450**, the method **1400** terminates. At least in some embodiments, the estimation rules and ranking rules may be tailored to suit the needs of a particular case.

Accordingly, at least some embodiments disclosed herein provide techniques for evaluating an application for deployment and detecting and resolving deployment conflicts between applications. One embodiment provides a deploy-

23

ment tool that receives a request is received to execute the application. Upon determining that insufficient execution data is present for the application, the deployment tool deploys the application for execution on a general-purpose processing environment. Upon determining that sufficient 5 execution data is present for the application, the deployment tool deploys the application for execution partially in time on a special-purpose processing environment, based on the execution data.

In some embodiments, a deployment conflict may be 10 detected based on application-specific measures such as usage amount of resources of a given type, usage duration of resources of the given type, and a reference time window in which to determine conflicts. The deployment conflict may be resolved by modifying one or more execution plans of appli- 15 cations involved in the conflict, in order to redeploy at least one application to a different processing environment for at least part of the time window in which the deployment conflict is determined to occur. The deployment conflict may be resolved in favor of certain applications based on criteria such 20 as service level agreements of the applications, expected performance improvement of the applications, and expected resource usage amount and/or duration of the applications.

Consequently, the application may execute more effi- 25 ciently using—even if only partially in time—the enhanced capabilities provided by one or more special-purpose processing environments. Further, the user providing the application may adjust the service level agreement to manage the monetary cost incurred from executing the application on heterogeneous processing environments. Consequently, the 30 user need not endure performance associated with executing the application from start to finish on a general-purpose processing environment, which may be undesirably slow to the user. Further, the user need not incur the monetary cost associated with executing the application from start to finish on a 35 special-purpose processing environment, which may be prohibitively expensive for the user. The user also need not incur the effort involved in attempting to determine, with some degree of accuracy, special-purpose processing environment(s) suited to the resource demands of the application. 40 Further, processing environments with distinct capabilities in a heterogeneous computing environment may also be utilized more effectively, because applications are migrated to those processing environments most equipped to satisfy the resource demands of the application—as the resource 45 demands change over time during execution of the application. Further still, detecting and resolving deployment conflicts may allow the applications to execute more efficiently in the processing environments at least in some cases.

While the foregoing is directed to embodiments presented in this disclosure, other and further embodiments may be devised without departing from the basic scope thereof, and the scope thereof is determined by the claims that follow.

What is claimed is:

1. A computer program product to manage deployment conflicts resulting from mid-execution application redeploy- ment based on activity patterns and performance metrics, the computer program product comprising:

a non-transitory computer-readable medium having pro- 60 gram code embodied therewith, the program code executable by one or more computer processors to:

prior to deploying a first application for execution:

determine an activity pattern based on one or more trial 65 executions of the first application on a first processing environment, the activity pattern characterizing resource demands of the first application over time

24

and indicating that the first application as being resource intensive in terms of a first resource type; and determine performance metrics based on one or more trial executions of the first application on a plurality of candidate processing environments each providing a higher hardware capability than the first processing environment in terms of the first resource type;

deploy the first application for execution in the first processing environment, responsive to a request to execute the first application;

redeploy the first application mid-execution from the first processing environment to a second processing environment selected from the plurality of candidate processing environments based on the performance metrics, wherein the first application is redeployed at a point in time determined based on the activity pattern; and

subsequent to redeploying the first application mid-ex- 50 ecution, resolve a deployment conflict characterized by the first application and a second application concurrently competing for resources of the first resource type on the second processing environment during a first time interval in which the resources of the first resource type are of an amount insufficient to satisfy a collective demand of a plurality of applications including the first and second applications.

2. The computer program product of claim 1, wherein the point in time is measured relative to a time of commencing execution of the first application in the first processing environment, wherein redeploying the first application mid-ex- 55 ecution includes suspending execution of the first application on the first processing environment and resuming execution of the first application on the second processing environment, wherein execution is resumed from where the first application left off on the first processing environment, wherein deploy- ing the first application for execution comprises deploying the first application for non-trial execution, wherein the first application is redeployed amid the non-trial execution, wherein resolving the deployment conflict comprises:

modifying an execution plan of at least one of the plurality 60 of applications such that the at least one application is redeployed for execution on a processing environment other than the second processing environment, for at least a portion of the first time interval, and such that the amount of the resources of the first resource type are sufficient to satisfy collective demand of the plurality of applications as a result of redeploying the at least one application.

3. The computer program product of claim 2, wherein the at 65 least one application is selected based on a set of predefined conflict resolution criteria, wherein the set of predefined conflict resolution criteria includes service level agreements of each of the plurality of applications, an expected performance improvement of each of the plurality of applications from using the resources of the first resource type in the second processing environment, an expected duration of each of the plurality of applications in using the resources of the first resource type in the second processing environment, and an expected amount of resources of the first resource type used by each of the plurality of applications.

4. The computer program product of claim 3, wherein the deployment conflict is resolved by a scheduler component of a predefined deployment tool, wherein the scheduler compo- 70 nent is configured to independently select the at least one application and modify the execution plan, based on each individual ranking selected from: an overall ranking of priority of the applications as determined based on predefined

25

application ranking criteria and a plurality of processing-environment-specific rankings of priority of the applications as determined based on the predefined application ranking criteria and processing environment ranking criteria.

5 5. The computer program product of claim 4, wherein the scheduler component is configured to independently resolve the deployment conflict between the first application and the second application at each individual point in time selected from: (i) prior to execution of the first application on the second processing environment; and (ii) during execution of the first application on the second processing environment, in which case any demand for the resources of the first resource type by the first application prior to the point in time is disregarded.

10 6. The computer program product of claim 5, wherein the first processing environment comprises a general-purpose processing environment, wherein the second processing environment comprises a special-purpose processing environment, wherein the scheduler component is configured to independently resolve the deployment conflict pertaining to each individual resource type selected from processor, memory, storage, network, graphics, and audio.

15 7. The computer program product of claim 1, wherein resolving the deployment conflict comprises:

modifying an execution plan of at least one application of the plurality of applications, wherein the at least one application is selected based on a service level agreement.

20 8. The computer program product of claim 1, wherein resolving the deployment conflict comprises:

modifying an execution plan of at least one application of the plurality of applications, wherein the at least one application is selected based on an expected performance improvement from using the resources of the first resource type.

25 9. The computer program product of claim 1, wherein resolving the deployment conflict comprises:

modifying an execution plan of at least one application of the plurality of applications, wherein the at least one application is selected based on an expected duration in using the resources of the first resource type.

30 10. The computer program product of claim 1, wherein resolving the deployment conflict comprises:

modifying an execution plan of at least one application of the plurality of applications, wherein the at least one application is selected based on an expected amount of resources of the first resource type used.

35 11. A system to manage deployment conflicts resulting from mid-execution application redeployment based on activity patterns and performance metrics, the system comprising:

one or more computer processors;

a memory containing a program which, when executed by the one or more computer processors, performs an operation comprising:

prior to deploying a first application for execution:

determining an activity pattern based on one or more trial executions of the first application on a first processing environment, the activity pattern characterizing resource demands of the first application over time and indicating that the first application as being resource intensive in terms of a first resource type; and

determining performance metrics based on one or more trial executions of the first application on a plurality of candidate processing environments

26

each providing a higher hardware capability than the first processing environment in terms of the first resource type;

deploy the first application for execution in the first processing environment, responsive to a request to execute the first application;

redeploying the first application mid-execution from the first processing environment to a second processing environment selected from the plurality of candidate processing environments based on the performance metrics, wherein the first application is redeployed at a point in time determined based on the activity pattern; and

subsequent to redeploying the first application mid-execution, resolving a deployment conflict characterized by the first application and a second application concurrently competing for resources of the first resource type on the second processing environment during a first time interval in which the resources of the first resource type are of an amount insufficient to satisfy a collective demand of a plurality of applications including the first and second applications.

12. The system of claim 11, wherein redeploying the first application mid-execution includes suspending execution of the first application on the first processing environment and resuming execution of the first application on the second processing environment, wherein execution is resumed from where the first application left off on the first processing environment, wherein deploying the first application for execution comprises deploying the first application for non-trial execution, wherein the first application is redeployed amid the non-trial execution, wherein resolving the deployment conflict comprises:

modifying an execution plan of at least one of the plurality of applications such that the at least one application is redeployed for execution on a processing environment other than the second processing environment, for at least a portion of the first time interval, and such that the amount of the resources of the first resource type are sufficient to satisfy collective demand of the plurality of applications as a result of redeploying the at least one application.

13. The system of claim 12, wherein the at least one application is selected based on a set of predefined conflict resolution criteria, wherein the set of predefined conflict resolution criteria includes service level agreements of each of the plurality of applications, an expected performance improvement of each of the plurality of applications from using the resources of the first resource type in the second processing environment, an expected duration of each of the plurality of applications in using the resources of the first resource type in the second processing environment, and an expected amount of resources of the first resource type used by each of the plurality of applications.

14. The system of claim 13, wherein the deployment conflict is resolved by a scheduler component of a predefined deployment tool, wherein the scheduler component is configured to independently select the at least one application and modify the execution plan, based on each individual ranking selected from: an overall ranking of priority of the applications as determined based on predefined application ranking criteria and a plurality of processing-environment-specific rankings of priority of the applications as determined based on the predefined application ranking criteria and processing environment ranking criteria.

15. The system of claim 14, wherein the scheduler component is configured to independently resolve the deployment conflict between the first application and the second application at each individual point in time selected from: (i) prior to execution of the first application on the second processing

27

environment; and (ii) during execution of the first application on the second processing environment, in which case any demand for the resources of the first resource type by the first application prior to the point in time is disregarded.

16. The system of claim 15, wherein the higher maximum capability comprises a higher maximum hardware capability, wherein the first processing environment comprises a general-purpose processing environment, wherein the second processing environment comprises a special-purpose processing environment, wherein the scheduler component is configured to independently resolve the deployment conflict pertaining to each individual resource type selected from processor, memory, storage, network, graphics, and audio, wherein each processing environment comprises a respective, separate hardware environment.

17. The system of claim 11, wherein resolving the deployment conflict comprises:

modifying an execution plan of at least one application of the plurality of applications, wherein the at least one application is selected based on a service level agreement.

28

18. The system of claim 11, wherein resolving the deployment conflict comprises:

modifying an execution plan of at least one application of the plurality of applications, wherein the at least one application is selected based on an expected performance improvement from using the resources of the first resource type.

19. The system of claim 11, wherein resolving the deployment conflict comprises:

modifying an execution plan of at least one application of the plurality of applications, wherein the at least one application is selected based on an expected duration in using the resources of the first resource type.

20. The system of claim 11, wherein resolving the deployment conflict comprises:

modifying an execution plan of at least one application of the plurality of applications, wherein the at least one application is selected based on an expected amount of resources of the first resource type used.

* * * * *